

Modeling Multi-attribute Negotiations in the Navy Detailing Process¹

Guoming Lai

Cuihong Li

Katia Sycara

Joseph Giampapa

January 18, 2005

CMU-RI-TR-05-06

Carnegie Mellon University

Robotics Institute

Technical Report

¹ This work was supported by the Navy Personnel Research, Studies and Technology (NPRST) under Navy contract number N6610-98-D-9501

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 18 JAN 2005		2. REPORT TYPE		3. DATES COVERED 00-00-2005 to 00-00-2005	
4. TITLE AND SUBTITLE Modeling Multi-attribute Negotiations in the Navy Detailing Process				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University,Robotics Institute,Pittsburgh,PA,15213				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT In this report, we first discuss the modeling context that is suitable for the multi-attribute negotiations in the Navy detailing process. Then, we present the model based on a novel negotiation protocol for automated multi-attribute negotiations. The negotiation algorithm, negotiation strategies and the mediator's problem are included in this model. Based on the discussions in this report, we conclude that our model is distributed, Paretoefficient and tractable.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 23	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Abstract

In this report, we first discuss the modeling context that is suitable for the multi-attribute negotiations in the Navy detailing process. Then, we present the model based on a novel negotiation protocol for automated multi-attribute negotiations. The negotiation algorithm, negotiation strategies and the mediator's problem are included in this model. Based on the discussions in this report, we conclude that our model is distributed, Pareto-efficient and tractable.

1 Introduction

Multi-attribute negotiation is a negotiation that involves multiple issues that need to be negotiated simultaneously. Usually it is characterized by two agents who have common interests as well as conflicts on the issues and want to seek an agreement about how to divide the total payoffs.

Multi-attribute negotiation is a useful mechanism in real life, not only because there are situations where people must deal with multiple issues simultaneously, such as an item-fixed supply contract on quantities, price and time, but also because people often are willing to introduce additional issues into their negotiations. A direct reason that people prefer to negotiate multiple issues simultaneously is that both sides can achieve benefits through the multi-attribute negotiations. For example, when selling automobiles agents usually want to add insurance, warranty and etc. into the contract rather than the single price of the automobile. This is because by having some discount on the insurance and warranty, buyers may accept the automobile price. On the other hand, buyers also find it to their interest to negotiate such a package because the price of buying insurance individually may be much higher. Such situations are normally called “Win-Win” situations.

Multi-attribute negotiations are very important for the Navy detailing process [1]. Most of the contracts between commands and sailors must include multiple issues, such as payment rate, projected rotation date, length of service, training, etc. In such situations, commands and sailors negotiate and reach agreement on all those issues. However, by negotiating multiple issues, commands and sailors can also be better off for the following reasons: As commands and sailors may have different preferences on issues, multi-attribute negotiations can enlarge the agreement zone, by trading-off on different issues. Thus, multi-attribute negotiations can increase the chance to reach good agreements. In addition, since commands and sailors may have different preferences in multi-attribute negotiations, they also may achieve agreement of high utility by trading of concessions on less important attributes for gains in more important attributes. We know these benefits are not available in a single-attribute negotiation because it is a “Win-Lose” situation [1] [2].

However, it is not trivial to implement multi-attribute negotiations and realize the benefits, because a multi-attribute negotiation is much more complex than a single-attribute negotiation [1]. In [1] we see there exist some multi-attribute negotiation models. However, these models are not practical in real-world applications like the Navy detailing process because most of them make strict assumptions, for example, complete information, explicit and linear additive utility functions, and independent reservation price on each issue. Although these assumptions can simplify a multi-attribute negotiation problem, they are not realistic in most real-world application domains. Thus, in this report we need to develop a novel model for multi-attribute negotiations based on the following principles:

- Fidelity: The model should reflect most of the main features and concerns of multi-attribute negotiations in the Navy detailing process.
- Tractability: The model should incorporate assumptions that make the quantitative description and analysis tractable without sacrificing the fidelity of the model.
- Generality: The model should be general. It also can be applied to other real-world applications besides the Navy detailing process, as long as the assumptions can be satisfied.

When building a model which is capable of flexible and sophisticated multi-attribute negotiations in realistic situations such as the Navy detailing process, the first thing that must be considered is the modeling context. By modeling context, we mean the characteristics of the environment in which the negotiations arise. Generally, the modeling context can include: assumptions about information, agents' preferences and utility functions, the relationship between agents, and the availability of intermediaries. The modeling context can impact the characteristics and the applicability of a model.

For multi-attribute negotiations, normally we can divide the modeling context by information setting and whether agents' utility functions are known. Thus, there can be four types of modeling contexts for multi-attribute negotiations: (i) complete information & utility function known; (ii) incomplete information & utility function known; (iii) complete information & utility function unknown; and (iv) incomplete information & utility function unknown. By complete or incomplete information, we mean whether

agents can have private information on their preferences, utility functions and negotiation strategies. By utility function known or unknown, we mean whether agents have explicit utility functions that represent their preferences on the issues. In particular, the situation that “utility function is unknown” is that agents don’t have explicit mathematical functions that express their preferences on the n -dimensional negotiation space, although they could tell their preferences on some given points in that space. In this report, we concentrate on the context of the fourth type with incomplete information & utility function unknown. We know this context is the one closest to realistic situations and the one where there is hardly any research work to date.

A model for automated multi-attribute negotiations, needs to include negotiation protocol, negotiation algorithm and negotiation strategies. To this end, the model presented in this report defines a novel negotiation protocol for multi-attribute negotiations and specifies the key structures and processes involved in this endeavor. We present a formal algorithm and heuristic strategies for agents, which can be implemented in realistic situations. The main contributions of this model are: (i) it is shaped by practical considerations as well as strictly theoretical insights; (ii) it defines a novel protocol for multi-attribute negotiations; (iii) it presents a formal algorithm and heuristic negotiation strategies, and (iv) it can be implemented in realistic situations. We say the model we propose is distributed, tractable, effective and efficient².

The rest of the paper is organized as follows: in Section 2 we first discuss the modeling context of our research, and then we present the formal model for the Navy detailing process. Section 3 gives the conclusions and discusses on future work.

2 The model

2.1 The modeling context

The goal of this report is to build a distributed, tractable, effective and efficient model for multi-attribute negotiations so that it can be applied in the Navy detailing process, as well

² By “distributed”, we mean the model considers agents as self-interested players and agents make their decisions by themselves in the negotiation. By “tractable and effective”, we mean the model can work under realistic environments and converges to an outcome for a multi-attribute negotiation problem effectively, no matter whether there is an agreement zone or not. By “efficient”, we mean our model will reach a Pareto-efficient agreement if there is an agreement zone for the multi-attribute negotiation problem.

as other real-world applications. In the Navy detailing process, we know commands and sailors usually don't have explicit utility functions that can represent their preferences on the whole negotiation space and the information setting in their negotiations is incomplete. So we say the modeling context we consider where information is incomplete and utility functions unknown fits well with the real situations in the Navy detailing process.

We briefly mention related work in simpler negotiation contexts. The simplest context for multi-attribute negotiations is the one with complete information and in which agents' utility functions are known. For this context, as agents know the utility functions (mathematical expressions) of each other, it is not hard to compute the Pareto-optimal frontier of their negotiation. So, rational agents can reach agreement on this frontier by Nash axioms or Rubinstein's alternating-offer game [4] [5] [7]. For the context where information is incomplete but utility functions are known, which means agents know their own indifference curves³ and their conceding directions on the negotiation space, agents may apply heuristic methods to negotiate, for instance, agents propose their indifference curves to the opponent by time-dependent strategy during each period. Another context besides ours is the one with complete information but utility functions are unknown. Although this context is close to our modeling context, it is still much simpler but less practical. First, because of the assumption of complete information, agents know the opponent's preference although they have no explicit utility functions. So the only obstacle left is how to find the Pareto-frontier without knowledge of utility functions. Ehtamo's methodology [15] [16] can work in this context. Second, the work in this context is not practical because complete information is not common in real-world negotiations. By contrast, in our modeling context, an agent can have private information so that she can play any rational strategy that is beneficial for herself to compete with her opponent without letting her opponent know how much she already conceded and how much she still can concede.

We, however, need some reasonable assumptions to make our model tractable:

Assumption 1: We assume the preference of each agent is *quasi-concave*. For any solution x , if the set of solutions that an agent prefers to x is convex, then the preference of this

³ An indifference curve is the one on which all the points are indifferent for an agent.

agent is *quasi-concave* (see Figure 1). It implies if there exist solutions that are better than x , then any convex combination (1-sum weight linear combination, e.g. $\lambda_1 x_1 + \lambda_2 x_2$, where $\lambda_1 + \lambda_2 = 1$ and $\lambda_1, \lambda_2 \geq 0$) of these solutions is still better than x . For example, if x' and x'' in Figure 1 are better than x for agent1, then any solution on the line connecting x' and x'' is also better than x . This assumption is reasonable in realistic situations and fits people's preference well. In contrast, most of the existing research on multi-attribute negotiations assumes agents have linear and additive utility functions, which is a much stricter assumption. The benefit of this “*quasi-concave*” assumption is that it implies the indifference curves⁴ of agents are one-peak so that each Pareto-efficient solution of a multi-attribute negotiation is on a joint tangent hyperplane of a pair of indifference curves of the two agents, and this hyperplane goes through this solution point. For instance, Y is on the joint tangent L_1 of the indifference curves C_1 and C_2 , which is a Pareto-optimal solution of the negotiation represented in Figure 1, and L_1 crosses Y .




Figure 1 Quasi-concave preference (Indifferent Curves and Pareto-optimal solution)

Assumption2: We assume agents know their most preferred points on the feasible solution space. This assumption is reasonable even for a multi-attribute negotiation. For example, the dashed rectangle Z in Figure 2 is the feasible (negotiable) negotiation space. In this space, it is reasonable to assume NI (say a command) knows that her most preferred

⁴ Although in our context agents' don't know their utility functions explicitly, the indifferent curves and conceding directions of their preferences do exist on the negotiation space.

point is “ x ”, and the most preferred point is “ y ” for $N2$ (say a sailor). Besides, we assume agents have enough computation/reasoning capacity such that they can tell their preferences on a given set of points on the negotiation space.

Assumption3: We assume a non-biased mediator is available. This assumption is also reasonable in the Navy detailing process since we know the detailer is playing such a role in the current system. In our model, the mediator’s responsibility is to help two agents maintain a Pareto-optimal solution.




Figure 2 The protocol of the model

2.2 The protocol of the model

Since in our context agents don’t have explicit utility functions and the information setting is incomplete, it is intractable to compute the whole Pareto-frontier of a multi-attribute negotiation problem. Besides, the existing methods to compute the Pareto-optimal frontier usually assume agents cooperate with each other and don’t consider agents’ rationality [15] [16]. Thus, we say these methods are not directly practical in our context where agents are assumed to be self-interested players. Furthermore, the existing models such as [17] [19] [20] [21] for multi-attribute negotiations that we introduce in

our literature review are also intractable in this context, because the protocols they apply usually need to search the whole negotiation space and reason on the opponent's preference, which can lead to a very heavy burden of computation and reasoning. Those models don't work well for an additional reason. They need the assumption of explicit utility functions. However, it is well known that preference elicitation is a time consuming process (or even intractable in some situations) [1]. In addition, inefficiency is another problem of most existing models as they seldom consider Pareto-efficiency of the negotiation result.

Faced with such problems, we design a model with a novel negotiation protocol. The protocol is explained as follows (via figures):

We assume agents know the most preferred points on the negotiation space. Therefore we can connect these two global optimal points “ x ” and “ y ” by a line L_I (see Figure 2). Then the first mover, say $N1$, makes the first proposal—a reference point “ m ” on L_I by her negotiation strategy (“ m ” should be a satisfying point for $N1$ because in this protocol it doesn't prohibit the responder from accepting “ m ” as the agreement of the negotiation.). Receiving $N1$'s proposal, $N2$ can accept this reference point “ m ” or reject it by her own negotiation strategy. If $N2$ accepts it, then the mediator works with $N1$ and $N2$ together to find the Pareto-optimal solution based on this reference point. When they reach a Pareto-optimal solution, say “ n ”, $N2$ still can reject this solution but $N1$ can't because it is $N1$ who made this proposal. (Actually, for both agents, the Pareto-optimal solution can not be worse than the reference point based on which it is reached, because in the searching procedure agents can always compare the candidate with this reference point. So if the candidate is worse than the reference point, they won't accept it.) So if $N2$ accepts “ n ” as the solution of the negotiation, then the negotiation ends with this outcome, and we say that this outcome is confirmed by both sides based on their own negotiation strategies. Thus, it is a decentralized solution. Moreover, the solution arrived at in this way is Pareto-optimal. If $N2$ rejects “ m ” or “ n ”, then it is the turn of $N2$ to propose an alternate offer—another reference point. (In the scope of this discussion, we assume that the searching cost is negligible. But we need to mention that in the situations where time is very valuable and the searching cost for agents to find a Pareto-optimal point is not negligible, agents also may accept the reference point at once as the outcome of the

negotiation. In such situations, the outcome may not be Pareto-optimal. However, we will show the effectiveness of our protocol, and in addition, we also will show that this protocol can make the reference point closer and closer to a Pareto-optimal point. Thus, even when the time is very valuable, the final reference point can be close to a Pareto-optimal point. This argument will be verified by simulations in the future.)

Subsequently, “ n ” and “ y ” are connected by a line L_2 . Then $N2$ can make a proposal—a reference point on L_2 to $N1$ and the mediator, for example, “ p ” in Figure 3. Assume $N1$ accepts this reference point “ p ”. Then the mediator works with $N1$ and $N2$, and reaches a Pareto-optimal solution “ q ”. If $N1$ rejects this solution, then it is again $N1$ ’s turn to make a proposal.




Figure 3 The protocol of the model

Now “ n ” and “ q ” are connected by line L_3 , and $N1$ makes a proposal on L_3 (see Figure 4). They repeat this procedure till they reach an agreement or the negotiation breaks down. We call the line L_1, L_2, \dots, L_i etc. as the negotiation base line. ■

For rational agents whose goal is to maximize their own utility, they should not leave extra money on the table. Thus, if a negotiation model is said to be efficient, it should be

a model based on which rational agents can reach Pareto-optimal agreements if there exists an agreement zone. In the above protocol, agents first propose reference points to each other, and then if a reference point is accepted, they will work together to search a Pareto-optimal point based on this reference point. The procedure ends only if a Pareto-optimal agreement is reached or the negotiation breaks down. Thus, the goal of this protocol determines that the outcome of our model is Pareto-optimal.




Figure 4 The protocol of the model

Next, in this protocol, agents make decisions by themselves. First, the reference point they propose is generated based on their own utility considerations. They will not propose a point they are not satisfied with, because by proposing a reference point, the agent must confirm no matter if the opponent accepts the Pareto-optimal point reached based on this reference point or the reference point itself (in the special cases where time is valuable) as the outcome of the negotiation. To propose an unsatisfying reference point is not beneficial for agents. Second, to accept a reference point or a Pareto-optimal solution is also decided by agents themselves. For example, agents can reject any reference point and Pareto-optimal solution if they are not satisfied with them. Third, agents reach the decision independently whether to cooperate with the opponent. Because the outcome of

the procedure to search a Pareto-optimal solution is no worse than the reference point, it is always beneficial for the proposer to cooperate with her opponent. On the other hand, as in this protocol the responder still can reject the Pareto-optimal solution reached by their cooperation, it is no harmful for the responder to cooperate (assuming time is not very valuable). So the decisions on the cooperation are reached by their own utility considerations. Hence, this protocol is distributed.

Tractability and effectiveness are other two important issues for multi-attribute negotiations. A multi-attribute negotiation is much harder than a single-attribute negotiation because the negotiation space of a multi-attribute negotiation is *n-dimensional*. As a result, people faced with a multi-attribute negotiation problem usually don't know how to negotiate, which offer to propose, which offer to accept, which direction to concede, how much to concede, etc. Many researchers have proposed various searching methods such as [19] [20] [21] to negotiate multiple issues. But the problem of those methods is that the computation is complex and they are not tractable in real situations. We know most points on the negotiation space are not candidates for the negotiation outcome, or a set of indifferent points on the space can be represented by one point. In other words, it is not necessary to search the whole negotiation space. In the above protocol, a novel idea we have proposed is to reduce an *n-dimensional* negotiation to a *single-line* negotiation, where agents propose offers based on some negotiation base lines. The trick of this reduction is that it doesn't change the solution region (the set of Pareto-optimal solutions) of the negotiation problem, because the only candidates for the outcome of the negotiation are Pareto-optimal points and this protocol doesn't lose any Pareto-optimal point. Faced with the new problem, it is not hard any longer for agents (even for person negotiator) to decide how much to propose or how much to concede since the decision space becomes a line. Now the remaining issue for the protocol is how to tractably and effectively search Pareto-optimal solutions.

It is still not trivial to reach a Pareto-optimal solution in a multi-attribute negotiation even by the decomposition above. In [1], we introduced three existing methods which consider Pareto-efficiency in multi-attribute negotiations. They are Nash solution, Ehtamo's constraint proposal method and Faratin's similarity criteria method. For Nash solution [4] [5], we know it needs the assumptions of cooperation, complete information and explicit

utility functions. So it is not practical in our context. For Faratin's method [19], although the trading-off mechanism based on similarity criteria can improve the speed of convergence in a multi-attribute negotiation, Pareto-efficiency can not be necessarily maintained. That mechanism of agents making most similar offers to the opponent's previous offer may not necessarily lead to a solution that is on a joint tangent hyperplane of a pair of indifference curves of the two agents. The solution can be much worse than a Pareto-optimal solution. The problem for Ehtamo's method [15] [16] is that it doesn't consider agents as self-interested rational agents, so it is still a centralized method. Besides, the burden of computation is high in Ehtamo's method as it needs to do iterations in every period in order to find the Pareto-frontier. However, Ehtamo's method is the one that can maintain Pareto-efficiency in most of cases.

Thus, in our protocol we apply the idea of Ehtamo's method to maintain Pareto-efficiency. However, as discussed above, agents in the protocol propose and accept reference points by themselves. By doing so, it is ensured that the negotiation problem is a decentralized problem, in contrast to the centralized protocol of Ehtamo's, where the mediator chooses all reference points. What is more, in the protocol that an agent can reject reference points by her own negotiation strategy makes it not necessary to search Pareto-optimal solutions in all periods. This avoids doing iterations all the time and saves computational cost. Furthermore, the most important in our protocol is that we update the negotiation base line each time when a Pareto-optimal solution is found. This mechanism can make the negotiation base line closer and closer to the Pareto-frontier (see Figure 5). This mechanism can accelerate the speed to converge to the result of the negotiation because it eliminates all the candidates that are worse than the current Pareto-optimal point for the opponent. We can imagine that if the opponent rejects the current Pareto-optimal point, she surely will reject the points worse than this point (An assumption here is that agents are rational and will not accept the offers which she ever rejected). Thus, it is not necessary to reach those points any longer. Another point of this mechanism is that it may improve the speed to find a Pareto-optimal point since the reference point is approaching to the Pareto-frontier during the procedure (This property will be rigorously verified by simulations). Besides, the property that the reference point turns to be closer and closer to a Pareto-optimal point can help to improve the efficiency of the outcome in

the situations where time is very valuable. Hence, this protocol improves Ehtamo's method and it behaves as a distributed, tractable, and effective method.

2.3 Negotiation algorithm

Thus, we can design a formal algorithm based on the protocol as follows:

Step1: The proposer in current period makes a proposal—a reference point on the current negotiation base line to the opponent and mediator by her negotiation strategy.

Step2: The responder makes her decision to accept/reject this reference point by her negotiation strategy. If the responder rejects it and the negotiation deadline is reached, then the algorithm ends with negotiation breakdown; otherwise, the two agents interchange their roles and the algorithm goes back to *step1*. If the responder accepts this reference point then the algorithm goes to *step3*.




Figure 5 The update mechanism for the negotiation base line

Step3: The mediator works with the two agents to find the Pareto-optimal solution based on the reference point in *step2*. The algorithm goes to *step4* when the solution is reached.

Step4: The responder makes her decision whether to accept/reject the Pareto-optimal solution found in *step3* by her negotiation strategy. If she accepts it then the algorithm ends with this solution; else if she rejects it but the deadline is reached, the algorithm also ends with the breakdown of the negotiation; otherwise, the two agents interchange their roles, and the algorithm updates the negotiation base line by connecting the current Pareto-optimal solution and the previous one (at the beginning, the previous one might be the global optimal point of an agent), and then the algorithm goes back to *step1*.

2.4 Negotiation strategy

The last two issues that need to be considered in our model are the negotiation strategies for the agents and the mediator's problem. Faced with such a complex context, we propose to consider appropriate heuristic strategies for agents in this model. There could be several suitable heuristics. We mention three as follows:

Heuristic 1: time-dependent strategy

As agents might know their global optimal point and global worst point⁵ on the given negotiation space, it is reasonable to assume the corresponding utility range of agents is $[0,1]$ with the bounds representing the global worst and global optimal points. So suppose x_w and x_o are the global worst and optimal points for $N1$, and y_w and y_o are the global worst and optimal points for $N2$ in the previous example. Then we have $U_1(x_w)=0$ and $U_1(x_o)=1$, and $U_2(y_w)=0$ and $U_2(y_o)=1$, where $U_1(\cdot)$ and $U_2(\cdot)$ are the utility functions for agents $N1$ and $N2$.

The agents can apply time-dependent strategy as equation 1:

$$U_i(t) = 1 - \left(\frac{t}{T}\right)^{1/\beta_i} \quad (1)$$

to decide how much utility to concede in each period, where T is the deadline, t is the current period and β_i is the strategy parameter of agent i . An agent can find the point on the current negotiation base line using her preference, which coincides with $U_i(t)$ in equation 1.

$$x_t = U_1^{-1}(t) \quad (2)$$

⁵ Similar as assumption2, it is also reasonable to assume agents know the worst solution for themselves on a given negotiation space.

$$y_t = U_2^{-1}(t) \quad (3)$$

Then if it is the turn of an agent to make a proposal, she can make the proposal characterized by this point. If an agent is a responder, she can compare the proposal from her opponent or the Pareto-optimal offer with this point and make the decision whether to accept/reject it. Because in our model agents are only faced with a single line in each period, we say it is not hard for computational agents to find out such points as in equation 2 and 3 on a line. As a result, this heuristic method is tractable.

Heuristic 2: fixed-distance conceding

In this heuristic method, an agent first divides the value range of her most important issue into several sub-ranges (the number of the sub-ranges can depend on the negotiation deadline). For instance, in Figure 6, agent *N1* applies this heuristic and divides the salary range [1000, 3000] into 5 sub-ranges with the division lines “ l_0 ” to “ l_4 ”. Then an agent makes proposals and decisions by the following strategies:

- **Proposal strategy:** At each turn, an agent proposes the point that is on the intersection between the current negotiation base line and a division line of the sub-ranges. For instance, “ x ”, “ m ” and “ r ” in Figure 6 might be the proposals agent *N1* will make. For convenience of explanation, we say “ x ” belongs to sub-range1, “ m ” is in sub-range2, and “ r ” is in sub-range3.
- **Decision strategy on a proposal (a reference point):** At each turn, an agent compares the current proposal with the sub-range in which she will make the proposal in the next period. If the current proposal is in that sub-range, then the agent accepts this proposal and works with the mediator and her opponent to find the corresponding Pareto-optimal point. Otherwise, she rejects it. For example, *N2* makes a proposal of the reference point “ p ” to *N1*, which is in sub-range4. But because the next proposal *N1* will make is on the division line “ l_2 ” that is in sub-range3, *N1* rejects this proposal (see Figure 6).
- **Decision strategy on a Pareto-optimal solution:** At each turn, an agent compares the Pareto-optimal solution with the proposal she will make in the next period. If the Pareto-optimal solution is better than her next proposal, the agent accepts this solution. Otherwise, she rejects it. For instance, we assume *N2* makes a proposal “ q ”

Heuristic 3: fixed-preference conceding

⁶ By equal division, we mean, for example, if NI can tell she prefers “ x ” (the global optimal point) 5 times to “ y ” (we assume it is the global worst point for NI), then she prefers “ x (d_0)” 1 time to “ d_1 ”, 2 times to “ d_2 ”, 3 times to “ d_3 ” and 4 times to “ d_4 ”. In application, of course agents also can apply unequal segment division.

- **Proposal strategy:** At each turn, an agent proposes the point on the current negotiation base line, which has the same preference as the corresponding division point on the first negotiation base line. For instance, in Figure 7, the first time, *N1* makes the proposal “ x (d_0)” on negotiation base line “ L_1 ”. In *N1*’s second turn, if the negotiation base line is still “ L_1 ”, then *N1* makes the proposal “ d_1 ”; else if the negotiation base line is “ L_2 ”, then *N1* tries to find out the point on “ L_2 ” which has equal preference as “ d_1 ” (i.e. “ m ” in the figure), and proposes it to *N2*. Similarly, “ r ” might be *N1*’s third proposal on negotiation base line “ L_3 ”, which has equal preference as “ d_2 ”.
- **Decision strategy on a proposal (a reference point):** Here, we assume “ d_i ” is division point based on which the decision agent will make her alternate proposal if she rejects the current proposal now. The decision agent compares the current proposal with “ d_{i+1} ”⁷. If the proposal is better than “ d_{i+1} ”, she accepts it; otherwise, she rejects it and then makes an alternate proposal based on “ d_i ”. For example, “ p ” is the proposal *N2* made to *N1*, but for *N1* it is worse than “ $d_{3(=2+1)}$ ” (assume “ d_2 ” is the division point based on which *N1* will make an alternate proposal if she rejects “ p ”), so *N1* rejects “ p ”(see Figure 7).
- **Decision strategy on a Pareto-optimal solution:** At each turn, an agent compares the Pareto-optimal solution with the division point based on which she will make an alternate proposal if she rejects this solution now. If the Pareto-optimal solution is better than that division point, an agent accepts this solution. Otherwise, she rejects it. For instance, we assume *N2* makes a proposal “ q ” rather than “ p ” in the above example. Thus, *N1* now accepts this proposal “ q ”, and works with *N2* and the mediator to find out a Pareto-optimal solution. Assume it is “ s ”. But for *N1*, “ s ” is worse than “ d_2 ”, so *N1* rejects this Pareto-optimal solution and makes an alternate proposal “ r ” on “ L_3 ” (see Figure 7).

⁷ Here, we propose that agents compare the proposal with “ d_{i+1} ” rather than “ d_i ” to make decision. The justification is that by doing so agents may reach agreement earlier than by the method of comparing the proposal with “ d_i ”. We know the Pareto-optimal solution can be much better than the corresponding reference point. So when the opponent’s proposal is better than “ d_{i+1} ”, there might already exist a chance to reach agreement.

2.5 The mediator's problem

For the mediator's problem, we apply the idea of Ehtamo's constraint proposal method to find Pareto-optimal solutions [15] [16]. So, the procedure in each turn to find out a Pareto-optimal solution can be described as following:




Figure 7 Fixed-preference conceding

Step 1. The mediator chooses a constraint (a line or a plane) going through the reference point and announces it to agents.

Step 2. Agents choose their most preferred points under the constraint and feed back to mediator.

Step 3. If the two points coincide or are very close, a Pareto-optimal solution is found. The procedure terminates. Else, the procedure goes back to *Step 1*.

The mechanism to choose constraints can be fixed point iteration, Newton's iteration or quasi-Newton's iteration discussed in [15].

3 Conclusions and future work

In this report, we presented a formal model for multi-attribute negotiations, which includes the modeling context, negotiation protocol, negotiation algorithm, negotiation strategies and the mediator's mechanism. We say our model has the following properties:

First, we model the multi-attribute negotiations under a context that is closest to real-world applications.

Second, the assumptions of our model are reasonable and flexible.

Third, our model is tractable and effective. We say autonomous computer agents have enough computation capacity to deal with the negotiation algorithm and strategies that we introduced above. Our model is also effective. The model appropriately decomposes the original multi-attribute negotiation on the whole *n-dimensional* space to a multi-attribute negotiation on *single-dimensional* lines (the negotiation base lines). This can decrease the load of computation and reasoning. What's more, in our model, agents can reject reference points so that it is not necessary to compute iterations in every period. Moreover, we propose a method for updating the negotiation base line. This method can make the negotiation base line converge to the Pareto-frontier such that it can make the negotiation more effective and improve the procedure to reach a Pareto-optimal point.

Fourth, the negotiation agreement if reached based on our model is Pareto-efficient.

Fifth, our model sufficiently considers agents' rationality. In our model agents make proposals or acceptance decisions based on their own negotiation strategies. Hence, our model is decentralized.

However, the efficiency and practical properties still need to be verified by simulations. The framework developed in this report lays out the foundation for our work in the next stage to analyze the model and provide an integrative solution for the negotiation decision problem in the Navy detailing problem. Another issue that might need to be researched more is the mediator's problem, which is also an important part in our model. In this report, we apply Ehtamo's method directly to maintain Pareto-efficiency. However, we know the method might not be always applicable in different situations. Thus, we still need to analyze it and may make further improvements. Moreover, Ehtamo's model assumes the presence of an unbiased mediator. However, this

assumption may not hold in many realistic situations. Therefore, in future work, we may extend our model to one in which a mediator is not unbiased or eliminate the mediator altogether.

6 References

- [1] Literature review of multi-attribute negotiations, Guoming Lai, Cuihong Li, Katia Sycara and Joe Giampapa, Nov.6, 2004.
- [2] H. Raiffa. *The Art and Science of Negotiation*. Harvard University Press, Cambridge, MA, 1982.
- [3] Cognitive Agent Technologies Project: Knowledge Engineering Study. Resource Consultants, Inc., Vienna, 2003.
- [4] J. Nash. Two-person Cooperative Games. *Econometrica*, 21(1):128-140, 1953.
- [5] J. Nash. The Bargaining Problem. *Econometrica*, 18(2):155-162, 1950.
- [6] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.
- [7] A. Rubinstein. Perfect Equilibrium in a Bargaining Model. *Econometrica*, 50(1):97-109, 1982.
- [8] A. Rubinstein. A Bargaining Model under Incomplete Information. *Econometrica*, 53:1151-1172, 1985.
- [9] M. Bac and H. Raff. Issue-by-Issue Negotiations: the Role of Information and Time Preference. *Games and Economic Behavior*, 13:125-134, 1996.
- [10] L.-A. Busch and I. J. Horstmann. Signaling via an Agenda in Multi-Issue Bargaining with Incomplete Information. *Economic Theory*, 13:561-575, 1999.
- [11] K. Lang and R. W. Rosenthal. Bargaining Piecemeal or All at Once. *The Economic Journal*, 111 (July):526-540, 2001.
- [12] L.-A. Busch and I. J. Horstmann. Bargaining Frictions, Bargaining Procedures and Implied Costs in Multiple-Issue Bargaining. *Economica*, 64:669-680, 1997.

- [13] E. Kalai. Proportional Solutions to Bargaining Situations: Intertemporal Utility Comparisons. *Econometrica*, 45:1623-1630, 1977.
- [14] R. B. Myerson. Two-Person Bargaining Problems and Comparable Utility. *Econometrica*, 45:1631-1637, 1977.
- [15] H. Ehtamo, et al. Generating Pareto Solutions in a Two-Party Setting: Constraint Proposal Methods. *Management Science*, 45:1697-1709, 1999.
- [16] P. Heiskanen. Generating Pareto-Optimal Boundary Points in Multiparty Negotiations Using Constraint Proposal Method, Wiley, New York, 2001.
- [17] S. S. Fatima, M. J. Wooldridge and N.R. Jennings. An Agenda-Based Framework for Multi-Issue Negotiation. *Artificial Intelligence*, 152:1-45, 2004.
- [18] X. Luo, N. R. Jennings, N. Shadbolt, H. Leung and J. H. Lee. A Fuzzy Constraint Based Model for Bilateral Multi-Issue Negotiations in Semi-Competitive Environments. *Artificial Intelligence Journal*, 148(1-2):53-102, 2003.
- [19] P. Faratin, C. Sierra, N.R. Jennings. Using similarity Criteria to Make Negotiation Trade-Offs. In the Proceeding of the 4th International Conference on Multi-Agent Systems, 119-126, Boston, USA, 2000.
- [20] M. Klein, P. Faratin, H. Sayama and Y. Bar-Yam. Negotiating Complex Contracts. MIT Sloan Working Paper No. 4196-01, 2001.
- [21] C. Li and G. Tesauro. A strategic Decision Model for Multi-attribute Bilateral Negotiation with Alternating Offers. ACM Conference on Electronic Commerce 2003: 208-209, 2003.
- [22] R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Cambridge University Press, 1976.
- [23] T.L. Satty. *The Analytic Hierarchical Process*. McGraw-Hill International, New York, USA, 1980.
- [24] M. Ehrgott. *Multicriteria Optimization*. Berlin: Springer-Verlag, 2000.
- [25] R. Steuer. *Multiple Criteria Optimization: Theory, Computation and Application*. New York, John Wiley & Sons, Inc., 1986.

- [26] C. Li, J. Giampapa and K. Sycara. A Solution for Bilateral Negotiations in Navy Detailing Process. Technical Report, Carnegie Mellon University, 2003.
- [27] C. Li, J. Giampapa and K. Sycara. A Review of Research Literature on Bilateral Negotiations. Technical Report, Carnegie Mellon University, 2003.
- [28] Sycara, K. "Problem Restructuring in Negotiation", *Management Science*, Vol. 37, No. 10, October, 1991.
- [29] Sycara, K. "Persuasive Argumentation in Negotiation", *Theory and Decision*, Vol. 28, No. 3, pp. 203-242, May 1990.
- [30] Sycara, K. "Negotiation Planning: An AI Approach", *European Journal of Operational Research*, Vol. 46, No. 2, pp. 216-234, May 1990.
- [31] Sycara, K. "Utility Theory in Conflict Resolution", *Annals of Operations Research*, Vol. 12, pp. 65-84, 1988.